

CLAIMS

What is claimed is:

1. A distributed data repository node allowing for content-based data storage and retrieval,
5 comprising
 an interface operative to communicate with client nodes and at least one other
distributed data repository node over a computer network, wherein client nodes transmit
messages including message payloads;
 a content-based filter layer operative to
10 receive message payloads;
 apply at least one content-based filter to the message payloads to extract one or
more content attributes defined in the at least one content-based filter from the message
payloads;
 a mapping module comprising a content map and at least one index map corresponding
15 to an attribute; wherein the mapping module is operative to
 receive message payloads and content attribute values extracted by the content-
based filter layer and associated with the message payloads,
 store the message payloads in the content map in association with unique
identifiers,
20 store the content attribute values associated with the message payloads in a
corresponding index map;
 synchronize the content attribute values in the at least one index map with the
at least one other distributed data repository node.
- 25 2. The distributed data repository node of claim 1 wherein the at least one content based filter
includes at least one filter criterion specifying a content element and an extraction method
defining extraction of the content element from the message payloads.
3. The distributed data repository node of claim 1 wherein the content-based filter layer is
30 further operative to apply at least one policy to the message payloads to select one or more

content-based filters; and wherein application of the at least one content based filter is conditioned on selection of at least one content based filter.

4. The distributed data repository node of claim 2 wherein the content-based filter layer is
5 further operative to

receive a content-based filter including at least one content element and an extraction method defining extraction of the content element from the message payloads;

construct at least one index map based on the content elements and extraction methods defined in the content-based filter.

10

5. The distributed data repository node of claim 1 wherein the mapping module is further operative to transfer received message payloads to at least one other distributed repository node for replication.

15 6. The distributed data repository node of claim 5 wherein the mapping module is further operative to receive message payloads from at least one other distributed repository node for replication, and store the message payloads in the content map in association with the unique identifier assigned by the other distributed repository node.

20 7. The distributed data repository node of claim 1 further comprising

a persistent data store;

wherein the content map and the at least one index map are maintained in a volatile memory; and

wherein the mapping module is further operative to maintain the message payloads and
25 the at least one index map in the persistent data store.

8. The distributed data repository node of claim 1 wherein the content map and the at least one index map are maintained in a volatile memory; and further comprising

a persistence module comprising a persistent data store and a persistence management
30 module; wherein the persistence management module is operative to
receive message payloads and attributes from the mapping engine, and

store the message payloads and the attribute values in the persistent data store.

9. The distributed data repository node of claim 8 wherein the persistence management module is further operative to

5 purge the least recently used data object from the content map maintained in the volatile memory as required for new message payloads when the limits of the volatile memory are exceeded.

10. The distributed data repository node of claim 1 further comprising

10 an external messaging agent operative to establish and maintain connections with the at least one other distributed data repository node.

11. The distributed data repository node of claim 1 further comprising

15 a node controller operatively coupled with the network interface to interact with client nodes and that least one data repository node.

12. The distributed data repository node of claim 1 further comprising

20 an external messaging agent operatively coupled with the network interface to transmit and receive messages from the at least one other data repository node.

13. The distributed data repository node of claim 8 further comprising

25 a message queuing module operative to

 maintain at least one message queue,

and wherein the message queuing module is operative to

25 receive messages from the mapping engine and the persistence module,

 place the messages on appropriate message queues, and

 transmit the queued messages to the mapping engine and the persistence module.

14. The distributed repository node of claim 13 wherein the mapping engine and the persistence

30 module are operative to register with the message queuing module to receive messages from the at least one message queue.

15. The distributed repository node of claim 1 wherein the mapping engine is further operative to identify message payloads in the content map associated with at least one given attribute value; and

5 return the identified message payloads in response to a query including the at least one given attribute value.

16. The distributed data repository node of claim 1 wherein the content-based filter layer is operative to modify the query based on application of at least one content-based filter.

10

17. The distributed data repository node of claim 16 wherein the content-based filter layer is further operative to apply at least one policy to the query to select one or more content-based filters; and wherein application of the at least one content based filter to the query is conditioned on selection of at least one content based filter.

15

18. The distributed data repository node of claim 16 wherein the at least one content based filter includes at least one filter criterion specifying a content element and an extraction method defining extraction of the content element from the message payloads; and wherein the query is modified based on the at least one filter criterion.

20

19. The distributed data repository node of claim 1 further comprising a messaging server operative to maintain a directory of other distributed data repository nodes and client nodes to facilitate one-way messaging between the nodes.

25 20. The distributed data repository system of claim 1 wherein the message payloads contain at least one data object.

21. The distributed data repository system of claim 1 wherein the at least one data object is a BLOB.

30

22. A distributed data repository system, comprising

at least two distributed repository nodes, each distributed repository node comprising
a content map containing at least one message payload stored in association
with a message payload identifier, and
at least one index map containing at least one content attribute value and a
5 corresponding message payload identifier;
wherein each distributed repository node is operative to
receive message payloads;
apply at least one content-based filter to the message payloads to extract one or
more content attributes defined in the at least one content-based filter from the message
10 payloads;
store the message payloads in the content map in association with unique
identifiers,
store the extracted content attribute values associated with the message
payloads in a corresponding index map;
15 synchronize the content attribute values in the at least one index map with
other distributed data repository nodes.

23. The distributed data repository system of claim 22 wherein the at least one content based
filter includes at least one filter criterion specifying a content element and an extraction
20 method defining extraction of the content element from the message payloads.

24. The distributed data repository system of claim 22 wherein each distributed data
repository node is further operative to apply at least one policy to the message payloads to
select one or more content-based filters; and wherein application of the at least one content
25 based filter is conditioned on selection of at least one content based filter.

25. The distributed data repository system of claim 24 wherein each distributed data
repository node is further operative to
receive a content-based filter including at least one content element and an extraction
30 method defining extraction of the content element from the message payloads;

construct at least one index map based on the content elements and extraction methods defined in the content-based filter.

26. The distributed data repository system of claim 22 wherein each distributed data repository node further comprises a mapping engine operative to identify message payloads in the content map associated with at least one given content attribute value; and
return the identified message payloads in response to a query including the at least one given attribute value.

27. The distributed data repository system of claim 26 wherein each distributed data repository node is operative to modify the query based on application of at least one content-based filter.

28. The distributed data repository system of claim 27 wherein each distributed data repository node is further operative to apply at least one policy to the query to select one or more content-based filters; and wherein application of the at least one content based filter to the query is conditioned on selection of at least one content based filter.

29. The distributed data repository system of claim 27 wherein the at least one content based filter includes at least one filter criterion specifying a content element and an extraction method defining extraction of the content element from the message payloads; and wherein the query is modified based on the at least one filter criterion.

30. The distributed data repository system of claim 22 wherein each distributed data repository node is operative to exchange message payloads with other distributed repository nodes as needed to fulfill client requests.

31. The distributed data repository system of claim 22 further comprising
at least one client node, wherein the at least one client node comprises a data repository client module operative to maintain a messaging connection with at least one of the distributed data repository nodes.

32. The distributed data repository node of claim 22 wherein the at least one client node further comprises a client application operative to interact with the distributed data repository system through the data repository client module.

5 33. A data repository node, comprising

an interface operative to communicate with client nodes and at least one other distributed data repository node over a computer network, and

a mapping module comprising a content map and at least one index map corresponding to an attribute; wherein the mapping module is operative to

10 receive a request to insert a record from a first client node;

generate a unique identifiers in response to the record insertion request;

transmit the unique identifier to the client node;

receive an insertion message including at least one record attribute value;

15 store the at least one record attribute value associated with the insertion message in a corresponding index map in association with the unique identifier;

receive record chunks from the client node;

store the record chunks in the content map in association with the unique identifier;

20 receive a query from a second client node, wherein the record satisfies the query;

provide the record chunks associated with the unique identifier corresponding to the record;

provide addition record chunks to the second client node as they are received from the first client node.

25 34. The data repository node of claim 33 wherein the mapping engine is further operative to synchronize the record attribute values in the at least one index map with the at least one other distributed data repository nodes.

35. The data repository node of claim 34 wherein the mapping engine is further operative to
30 transmit the record chunks to at least one other data repository node for replication.

36. A distributed data repository system, comprising
at least two distributed repository nodes, each distributed repository node comprising
a content map containing at least one message payload stored in association
with a message payload identifier, and
5 at least one index map containing at least one content attribute value and a
corresponding message payload identifier;
wherein each distributed repository node is operative to
receive a request to insert a record from a first client node;
generate a unique identifiers in response to the record insertion request;
10 transmit the unique identifier to the client node;
receive an insertion message including at least one record attribute value;
store the at least one record attribute value associated with the insertion
message in a corresponding index map in association with the unique identifier;
receive record chunks from the client node;
15 store the record chunks in the content map in association with the unique
identifier;
receive a query from a second client node, wherein the record satisfies the query;
provide the record chunks associated with the unique identifier corresponding
to the record;
20 provide additional record chunks to the second client node as they are received
from the first client node;
synchronize the record attribute values in the at least one index map with the at
least one other distributed data repository nodes.
- 25 37. The data repository node of claim 36 wherein each distributed data repository node is
further operative to transmit the record chunks to at least one other data repository node for
replication.
38. The data repository node of claim 36 wherein each distributed data repository node is
30 further operative to request and receive from at least one other distributed data repository node
record chunks that match a query received from a client node.